arXiv:1510.01413v1 [cs.IT] 6 Oct 2015

BER ANALYSIS OF THE BOX RELAXATION FOR BPSK SIGNAL RECOVERY

Christos Thrampoulidis⋆, *Ehsan Abbasi*⋆, *Weiyu Xu*†, *Babak Hassibi*⋆

⋆Department of Electrical Engeeniring, Caltech, Pasadena, USA
†Department of ECE, University of Iowa

ABSTRACT

We study the problem of recovering an n-dimensional vector of $\{\pm 1\}^n$ (BPSK) signals from m noise corrupted measurements $\mathbf{y} = \mathbf{A}\mathbf{x}_0 + \mathbf{z}$. In particular, we consider the box relaxation method which relaxes the discrete set $\{\pm 1\}^n$ to the convex set $[-1, 1]^n$ to obtain a convex optimization algorithm followed by hard thresholding. When the noise $\mathbf{z}$ and measurement matrix $\mathbf{A}$ have iid standard normal entries, we obtain an exact expression for the bit-wise probability of error P_e in the limit of n and m growing and $\frac{m}{n}$ fixed. At high SNR our result shows that the P_e of box relaxation is within 3dB of the matched filter bound (MFB) for square systems, and that it approaches (MFB) as m grows large compared to n. Our results also indicates that as $m, n \to \infty$, for any fixed set of size k, the error events of the corresponding k bits in the box relaxation method are independent.

Index Terms— BER Analysis, Box Relaxation, BPSK, Matched Filter Bound, Maximum Likelihood Decoder

1. INTRODUCTION

The problem of recovering an unknown BPSK vector from a set of noise corrupted linearly related measurements arises in numerous applications, such as Massive MIMO [1, 2, 3, 4]. As a result, a large host of exact and heuristic optimization algorithms have been proposed. Exact algorithms, such as sphere decoding and its variants, become computationally prohibitive as the problem dimension grows. Heuristic algorithms such as zero-forcing, MMSE, decision-feedback, etc., [5, 6, 7] have inferior performances that are often difficult to precisely characterize.

One popular heuristic is the so called "Box Relaxation" which replaces the discrete set $\{\pm 1\}^n$ with the convex set $[-1, 1]^n$ [8]. This allows one to recover the signal via convex optimization followed by hard thresholding. Despite its popularity, very little is known about the performance of this method. In this paper, we exactly characterize its bit-wise error probability in the regime of large dimensions and under Gaussian assumptions.

The remainder of the paper is organized as follows. Some background and the formal problem definition are given in the rest of this section. The main result and a detailed discussion follows in Section 2. An outline of the proof is the subject of Section 3. Finally, the paper concludes in Section 4.

1.1. Setup

Our goal is to recover an n-dimensional BPSK vector $\mathbf{x}_0 \in \{\pm 1\}^n$ from the noisy multiple-input multiple output (MIMO) relation $\mathbf{y} = \mathbf{A}\mathbf{x}_0 + \mathbf{z} \in \mathbb{R}^m$, where $\mathbf{A} \in \mathbb{R}^{m\times n}$ is the MIMO channel matrix (assumed to be known) and $\mathbf{z} \in \mathbb{R}^m$ is the noise vector. We assume that $\mathbf{A}$ has entries iid $\mathcal{N}(0, 1/n)$ and $\mathbf{z}$ has entries iid $\mathcal{N}(0, \sigma^2)$. The normalization is such that the reciprocal of the noise variance σ^2 is equal to the Signal-to-Noise Ratio, i.e. $\mathrm{SNR} = 1/\sigma^2$.

The Maximum-Likelihood (ML) decoder. The ML decoder which maximizes the probability of error (assuming the $\mathbf{x}_{0,i}$ are equally likely) is given by $\min_{\mathbf{x}\in\{\pm 1\}^n} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2$. Solving the above, is often computationally intractable, especially when n is large, and therefore a variety of heuristics have been proposed (zero-forcing, mmse, decision-feedback, etc.) [9].

Box Relaxation Optimization. The heuristic we shall use, we refer to it as Box Relaxation Optimization (BRO). It consists of two steps. The first one involves solving a convex relaxation of the ML algorithm, where $\mathbf{x} \in \{\pm 1\}^n$ is relaxed to $\mathbf{x} \in [-1, 1]^n$. The output of the optimization is hard-thresholded in the second step to produce the final binary estimate. Formally, the algorithm outputs an estimate $\mathbf{x}^*$ of $\mathbf{x}_0$ given as

$$\hat{\mathbf{x}} = \arg \min_{-1\le \mathbf{x}_i \le 1} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2, \tag{1a}$$

$$\mathbf{x}^* = \mathrm{sign}(\hat{\mathbf{x}}), \tag{1b}$$

where the sign function returns the sign of its input and acts element-wise on input vectors.

Bit error probability. We evaluate the performance of the detection algorithm by the bit error probability P_e, defined as the expectation of the Bit Error Rate BER . Formally,

$$BER := \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}_{\{\mathbf{x}_i^* \neq \mathbf{x}_{0,i}\}}, \tag{2a}$$

$$P_e := \mathbb{E}\left[BER\right] = \frac{1}{n}\sum_{i=1}^{n} \Pr\left(\mathbf{x}_i^* \neq \mathbf{x}_{0,i}\right). \tag{2b}$$

2. MAIN RESULT

Our main result analyzes the P_e of the (BRO) in (1). We assume a large-system limit where $m, n \to \infty$ at a proportional rate δ. The SNR is assumed constant; in particular, it does not scale with n. Let $Q(\cdot)$ denote the Q-function associated with the standard normal density $p(h) = \frac{1}{\sqrt{2\pi}} \mathrm{e}^{-h^2/2}$.

Theorem 2.1 (P_e of the (BRO))**.** *Let P_e denote the bit error probability of the detection scheme in (1) for some fixed but unknown BPSK signal $\mathbf{x}_0 \in \{\pm 1\}^n$. For constant* SNR *and $\frac{m}{n} \to \delta \in (\frac{1}{2}, \infty)$, it holds:*

$$\lim_{n\to\infty} P_e = Q(1/\tau_*),$$

where τ_* is the unique solution to

$$\min_{\tau>0} \frac{\tau}{2}\left(\delta - \frac{1}{2}\right) + \frac{1/\mathrm{SNR}}{2\tau} - \frac{\tau}{2}\int_{\frac{2}{\tau}}^{\infty} \left(h + \frac{2}{\tau}\right)^2 p(h)\mathrm{d}h. \tag{3}$$

Theorem 2.1 derives a *precise* formula for the bit error probability of the (BRO). The formula involves solving a *convex* and deterministic minimization problem in (3). We outline the proof in Section 3. First, a few remarks are in place.

2.1. Remarks

Computing τ_*. It can be shown that the objective function of (3) is strictly convex when $\delta > \frac{1}{2}$. When $\delta < \frac{1}{2}$, it is well known that even the noiseless box relaxation fails [10]. (In fact, $\delta = \frac{1}{2}$ is the recovery threshold for this convex relaxation.) Thus, (3) has a unique solution τ_*. Observe that the problem parameters δ and SNR appear explicitly in (3); naturally then τ_* is indeed a function of those. The minimization in (3) can be efficiently solved numerically. In addition, owing to the strict convexity of the objective function, τ_* can be equivalently expressed as the unique solution to the corresponding first order optimality conditions.

Numerical illustration. Figure 1 illustrates the accuracy of the prediction of Theorem 2.1. Note that although the theorem requires $n \to \infty$, the prediction is already accurate for n ranging on a few hundreds.

Analysis of convex algorithms. We are able to predict the P_e of the detection algorithm in (1) by analyzing the performance of the convex algorithm in (1a). These type of convex algorithms, which minimize a least-squares function of the residual $\mathbf{y} - \mathbf{A}\mathbf{x}$ subject to a (typically non-smooth) constraint on $\mathbf{x}$, are often referred to in the (statistics and signal-processing) literature as LASSO-type algorithms. When the performance of such algorithms is measured in terms of the square-error $\|\hat{\mathbf{x}} - \mathbf{x}_0\|_2^2$, the recent line of works [11, 12, 13, 14] has led to precise results and a clear understanding of its asymptotic behavior. The analysis of these works builds upon the Convex Gaussian Min-max Theorem (CGMT) [14, Thm. 1], which is an extension of a classical Guassian comparison inequality due to Gordon [15]. Of interest to us is not calculating the squared-error of (1a) but rather the P_e. Thus, we manage to extend the precise analysis of the LASSO-type algorithm beyond the squared-error. To prove our result we require a slight generalization of the CGMT as it appears in [14].

P_e at high-SNR. It can be shown that when $\mathrm{SNR} \gg 1$, then $\tau_* = 1/\sqrt{(\delta - 1/2)SNR}$. This can be intuitively understood as follows: at high-SNR, we expect τ_* to be going to zero (correspondingly P_e to be small). When this is the case, the last term in (3) is negligible; then, τ_* is the solution to $\min_{\tau>0} \frac{\tau}{2}\left(\delta - \frac{1}{2}\right) + \frac{1/\mathrm{SNR}}{2\tau}$ which gives the derided result. Hence, for $\mathrm{SNR} \gg 1$,

$$\lim_{n\to\infty} P_e \approx Q(\sqrt{(\delta - 1/2)\cdot \mathrm{SNR}}). \tag{4}$$

In Figure 2 we have plotted this high-SNR expression for the $\log_{10}(P_e)$ vs its exact value as predicted by Theorem 2.1. It is interesting to observe that the former is actually a very good approximation to the latter even for small practical values of SNR. The range of SNR values for which the approximation is valid becomes larger with increasing δ. Heuristically, for $\delta > 0.7$ the expression in (4) is a good proxy for the true probability of error at practical SNR values.

Comparison to the matched filter bound. Theorem 2.1 gives us a handle on the P_e of (BRO) in (1) and therefore allows to evaluate its practical performance. Here, we compare the performance to an idealistic case, where all $n-1$, but 1, bits of $\mathbf{x}_0$ are known to us. As is customary in the field, we refer to the bit error probability of this case as the *matched filter bound* (MFB) and denote it by P_e^{MFB}. The (MFB) corresponds to the probability of error in detecting (say) $\mathbf{x}_{0,n} \in \{\pm 1\}$ from: $\tilde{\mathbf{y}} = \mathbf{x}_{0,n}\mathbf{a}_n + \mathbf{z}$, where $\tilde{\mathbf{y}} = \mathbf{y} - \sum_{i=1}^{n-1} \mathbf{x}_{0,i}\mathbf{a}_i$ is assumed known, and, $\mathbf{a}_i$ denotes the i^{th} column of $\mathbf{A}$. The ML estimate is just the sign of the projection of the vector $\tilde{\mathbf{y}}$ to the direction of $\mathbf{a}_n$. Without loss of generality assume that $\mathbf{x}_{0,n} = 1$. Then, the output of the matched filter becomes $\mathrm{sign}(\tilde{X})$, where $\tilde{X} = \|\mathbf{a}_n\|^2 + \sigma^2\nu$, where $\nu \sim \mathcal{N}(0,1)$. When $n \to \infty$, $\|\mathbf{a}_n\|^2 \xrightarrow{P} \delta$[1] . Hence, with probability one,

$$\lim_{n\to\infty} P_e^{MFB} = \lim_{n\to\infty} \mathbb{P}(\tilde{X} < 0) = Q(\sqrt{\delta \cdot \mathrm{SNR}}). \tag{5}$$

A direct comparison of (5) to (4) shows that at high-SNR, the performance of the (BRO) is $10\log_{10}\frac{\delta}{\delta - 1/2}$dB off that of the (MFB) . In particular, in the square case ($\delta = 1$), where the number of receive and transmit antennas are the same, the (BRO) is 3dB off the (MFB) . When the number of receive antennas is much larger, i.e. when $\delta \to \infty$, then the performance of the (BRO) approaches the (MFB) .

3. PROOF OUTLINE

For simplicity, we write $\|\cdot\|$ for the ℓ_2-norm.

The error vector. It is convenient to re-write (1a) by changing the variable to the error vector $\mathbf{w} := \mathbf{x} - \mathbf{x}_0$:

[1] We use $\xrightarrow{P}$ to denote convergence in probability with $n \to \infty$.

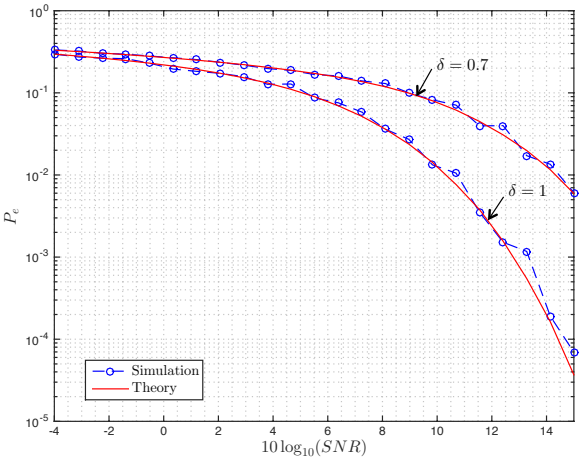

Fig. 1: BER Performance of the Boxed Relaxation: P_e as a function of SNR for different values of the ration $\delta = \lceil m/n \rceil$. The theoretical prediction follows from Theorem 2.1. For the simulations, we used $n = 512$. The data are averages over 20 independent realizations of the channel matrix and of the noise vector for each value of the SNR.

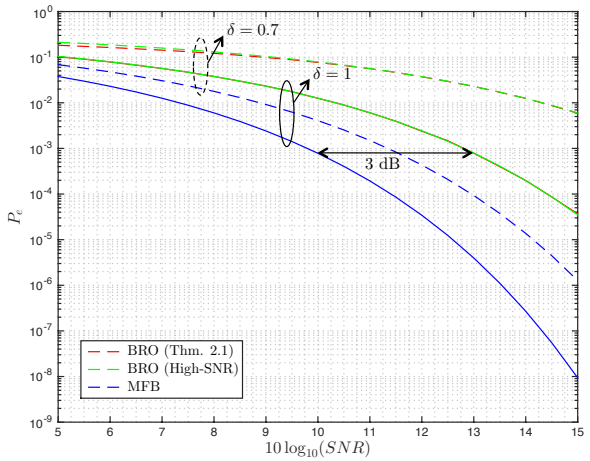

Fig. 2: Bit error probability of the Box Relaxation Optimization (BRO) in (1) in comparison to the Matched Filter Bound (MFB) for $\delta = 0.7$ (dashed lines) and $\delta = 1$ (solid lines). The red curves follow the formula of Thm. 2.1, the green ones correspond to (4), and, P_e^{MFB} of (5) is in blue.

$$\hat{\mathbf{w}} := \min_{-2\leq \mathbf{w}_i \leq 0} \|\mathbf{z} - \mathbf{A}\mathbf{w}\|. \tag{6}$$

Without loss of generality we assume for the analysis that $\mathbf{x}_0 = \mathbf{1}_n = (1, 1, \ldots, 1)$. Then, we can write (2a) in terms of the error vector $\mathbf{w}$ as: $BER = \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}_{\{\hat{\mathbf{w}}_i \leq -1\}}$.

The CGMT. The fundamental tool behind our analysis is the Convex Gaussian Min-max Theorem (CGMT) [14]; CGMT builds upon a classical result due to Gordon [15][2]. It associates with a primary optimization (PO) problem a simplified auxiliary optimization (AO) problem from which we can tightly infer properties of the original (PO), such as the optimal cost, the optimal solution, etc.. The idea of combining the GMT with convexity is attributed to Stojnic [11]. Thrampoulidis, Oymak and Hassibi built and significantly extended on this idea arriving at the CGMT as it appears in [14, Thm. 3] For the ease of reference we repeat a statement of the CGMT here; in this generality the theorem appears in [19]. Consider the following two min-max problems

$$\Phi(\mathbf{G}) := \min_{\mathbf{w}\in\mathcal{S}_\mathbf{w}} \max_{\mathbf{u}\in\mathcal{S}_\mathbf{u}} \mathbf{u}^T\mathbf{G}\mathbf{w} + \psi(\mathbf{w}, \mathbf{u}), \tag{7a}$$

$$\phi(\mathbf{g}, \mathbf{h}) := \min_{\mathbf{w}\in\mathcal{S}_\mathbf{w}} \max_{\mathbf{u}\in\mathcal{S}_\mathbf{u}} \|\mathbf{w}\|\mathbf{g}^T\mathbf{u} - \|\mathbf{u}\|\mathbf{h}^T\mathbf{w} + \psi(\mathbf{w}, \mathbf{u}), \tag{7b}$$

where $\mathbf{G} \in \mathbb{R}^{m\times n}, \mathbf{g} \in \mathbb{R}^m, \mathbf{h} \in \mathbb{R}^n$, $\mathcal{S}_\mathbf{w} \subset \mathbb{R}^n, \mathcal{S}_\mathbf{u} \subset \mathbb{R}^m$ and $\psi : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$. Denote $\mathbf{w}_\Phi := \mathbf{w}_\Phi(\mathbf{G})$ and $\mathbf{w}_\phi := \mathbf{w}_\phi(\mathbf{g}, \mathbf{h})$ any optimal minimizers in (7a) and (7b), respectively.

[2]Gordon's original result is classically used to establish non-asymptotic probabilistic lower bounds on the minimum singular value of Gaussian matrices (e.g. [16]), and has a number of other applications in high-dimensional convex geometry (e.g. [17, 18]).

Theorem 3.1 (CGMT). *In (7), let $\mathcal{S}_\mathbf{w}, \mathcal{S}_\mathbf{u}$ be convex and compact sets, ψ be continuous and convex-concave on $\mathcal{S}_\mathbf{w} \times \mathcal{S}_\mathbf{u}$, and, $\mathbf{G}, \mathbf{g}$ and $\mathbf{h}$ all have entries iid standard normal. Let $\mathcal{S}$ be an arbitrary open subset of $\mathcal{S}_\mathbf{w}$ and $\mathcal{S}^c = \mathcal{S}_\mathbf{w}/\mathcal{S}$. Denote $\phi_{\mathcal{S}^c}(\mathbf{g}, \mathbf{h})$ the optimal cost of the optimization in (7b), when the minimization over $\mathbf{w}$ is now constrained over $\mathbf{w} \in \mathcal{S}^c$. If in the limit of $n \to \infty$ both $\phi(\mathbf{g}, \mathbf{h})$ and $\phi_{\mathcal{S}^c}(\mathbf{g}, \mathbf{h})$ converge in probability, and, $\lim_{n\to\infty} \Pr(\mathbf{w}_\phi \in \mathcal{S}) = 1$, then, it also holds $\lim_{n\to\infty} \Pr(\mathbf{w}_\Phi \in \mathcal{S}) = 1$.*

Identifying the (PO) and the (AO). Using the CGMT for the analysis of the P_e, requires as a first step expressing the optimization in (1a) in the form of a (PO) as it appears in (7a). It is easy to see that (6) is equivalent to

$$\min_{-2\leq \mathbf{w}_i \leq 0} \max_{\|\mathbf{u}\|\leq 1} \mathbf{u}^T\mathbf{A}\mathbf{w} - \mathbf{u}^T\mathbf{z}. \tag{8}$$

Observe that the constraint sets above are both convex and compact; also, the objective function is convex in $\mathbf{w}$ and concave in $\mathbf{u}$. Hence, according to the CGMT we can perform the analysis of the BER for the corresponding (AO) problem instead, which becomes (note the normalization to account for the variance of the entries of $\mathbf{A}$)

$$\frac{1}{\sqrt{n}} \min_{-2\leq \mathbf{w}_i \leq 0} \max_{\|\mathbf{u}\|\leq 1} (\|\mathbf{w}\|\mathbf{g} - \sqrt{n}\mathbf{z})^T\mathbf{u} - \|\mathbf{u}\|\mathbf{h}^T\mathbf{w}. \tag{9}$$

We refer to the optimization in (9) as the (AO) problem.

Computing the BER via the (AO). Call $\tilde{\mathbf{w}}$ the optimal solution of the (AO). Fix any $\epsilon > 0$ and consider the set

$$\mathcal{S} = \{\mathbf{v} : \left|\frac{1}{n}\sum_{i=1}^{n} \mathbb{1}_{\{\mathbf{v}_i \leq -1\}} - Q(1/\tau_*)\right| < \epsilon\}, \tag{10}$$

where τ_* is defined in the statement of Theorem 2.1. We will apply Theorem 3.1 for the above set $\mathcal{S}$. In particular, we show

that (i) the (AO) in (9) converges in probability (after proper normalization with n), and, (ii) $\tilde{\mathbf{w}} \in \mathcal{S}$ with probability one. These will suffice to conclude that $\hat{\mathbf{w}} \in \mathcal{S}$ with probability one, which would complete the proof of Theorem 2.1.

Simplifying the (AO). We begin by simplifying the (AO) problem as it appears in (9). First, since both $\mathbf{g}$ and $\mathbf{z}$ have entries iid Gaussian, then, $\|\mathbf{w}\|\mathbf{g} - \sqrt{n}\mathbf{z}$ has entries iid $\mathcal{N}(0, \sqrt{\|\mathbf{w}\|^2 + n\sigma^2})$. Hence, for our purposes and using some abuse of notation so that $\mathbf{g}$ continues to denote a vector with iid standard normal entries, the first term in (9) can be treated as $\sqrt{\|\mathbf{w}\|^2 + n\sigma^2}\mathbf{g}^T\mathbf{u}$, instead. As a next step, fix the norm of $\mathbf{u}$ to say $\|\mathbf{u}\| = \beta$. Optimizing over its direction is now straightforward, and gives $\min_{-2\leq \mathbf{w}_i \leq 0} \max_{0\leq\beta\leq 1} \frac{\beta}{\sqrt{n}} \left(\sqrt{\|\mathbf{w}\|^2 + n\sigma^2}\|\mathbf{g}\| - \mathbf{h}^T\mathbf{w} \right)$. In fact, it is easy to now optimize over β as well; its optimal value is 1 if the term in the parenthesis is non-negative, and, is 0 otherwise. With this, the (AO) simplifies to the following:

$$\Big(\min_{-2\leq \mathbf{w}_i \leq 0} \sqrt{\frac{\|\mathbf{w}\|^2}{n} + \sigma^2}\|\mathbf{g}\| - \frac{1}{\sqrt{n}}\mathbf{h}^T\mathbf{w} \Big)_+,$$

where we defined $(\chi)_+ := \max\{\chi, 0\}$. To facilitate the optimization over $\mathbf{w}$, we express the term in the square-root in a variational form, using $\sqrt{\chi} = \min_{\tau>0} \frac{\tau}{2} + \frac{\chi}{2\tau}$. With this trick, the minimization over the entries of $\mathbf{w}$ becomes separable:

$$\min_{\tau\geq 0} \frac{\tau\|\mathbf{g}\|}{2} + \frac{\sigma^2\|\mathbf{g}\|}{2\tau} + \sum_{i=1}^{n} \min_{-2\leq \mathbf{w}_i \leq 0} \frac{\|\mathbf{g}\|}{2\tau n}\mathbf{w}_i^2 - \frac{\mathbf{h}_i}{\sqrt{n}}\mathbf{w}_i.$$

Then, the optimal $\tilde{\mathbf{w}}_i$ satisfies

$$\tilde{\mathbf{w}}_i = \begin{cases} 0 & , \text{if } \mathbf{h}_i \geq 0, \\ \frac{\tau\sqrt{n}}{\|\mathbf{g}\|}\mathbf{h}_i & , \text{if } -\frac{2\|\mathbf{g}\|}{\tau\sqrt{n}} \leq \mathbf{h}_i < 0, \\ -2 & , \text{if } \mathbf{h}_i < -\frac{2\|\mathbf{g}\|}{\tau\sqrt{n}}. \end{cases} \tag{11}$$

where, τ is the solution to the following:

$$\left(\min_{\tau>0} \frac{\tau\|\mathbf{g}\|}{2} + \frac{\sigma^2\|\mathbf{g}\|}{2\tau} + \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \upsilon(\tau; \mathbf{h}_i, \|\mathbf{g}\|) \right)_+, \tag{12}$$

$$\upsilon(\tau; \mathbf{h}_i, \|\mathbf{g}\|) := \begin{cases} 0 & , \text{if } \mathbf{h}_i \geq 0, \\ -\frac{\tau\sqrt{n}}{2\|\mathbf{g}\|}\mathbf{h}_i^2 & , \text{if } -\frac{2\|\mathbf{g}\|}{\tau\sqrt{n}} \leq \mathbf{h}_i < 0, \\ 2\frac{\|\mathbf{g}\|}{\tau\sqrt{n}} + 2\mathbf{h}_i & , \text{if } \mathbf{h}_i \leq -\frac{2\|\mathbf{g}\|}{\tau\sqrt{n}}. \end{cases}$$

Convergence of the (AO). Now that the (AO) is simplified as in (12), we can get a handle on the limiting behavior of the optimization itself as well as of the optimal $\tilde{\mathbf{w}}$. But first, we need to properly normalize the (AO) by dividing the objective in (12) by $\sqrt{n}$. Also, for convenience, redefine $\tau := \frac{\tau}{\sqrt{\delta}}$. By the WLLN, we have $\frac{\|\mathbf{g}\|}{\sqrt{n}} \xrightarrow{P} \sqrt{\delta}$, and, for all $\tau > 0$, $\frac{1}{n}\sum_{i=1}^{n} \upsilon(\tau; \mathbf{h}_i, \|\mathbf{g}\|) \xrightarrow{P} Y(\tau) := -\frac{\tau}{2}\int_0^{\frac{2}{\tau}} h^2 p(h)\mathrm{d}h + \frac{2}{\tau}Q\left(\frac{2}{\tau}\right) + 2\int_{\frac{2}{\tau}}^{\infty} hp(h)\mathrm{d}h$. With these we can evaluate the point-wise (in τ) limit of the objective function in (12). Next, we use the fact that the objective is convex in τ and Lemma [20, Cor.. II.1], to conclude that the convergence is indeed uniform in τ. Hence, the random optimization in (12) converges to the following deterministic optimization $\min_{\tau>0} \frac{\tau\delta}{2} + \frac{\sigma^2}{2\tau} + Y(\tau)$; some algebra shows that the latter is the same as (3). If $\delta > 1/2$, then, it can be shown via differentiation that the objective function of it is strictly convex. Also, it is nonnegative; thus, the entire expression in (12), which is nothing but the (AO) problem we started with, converges in probability to (3). What is more, using [21, Thm. 2.7] it can be shown that the optimal $\tau_*(\mathbf{g}, \mathbf{h})$ of the (AO) converges in probability to the unique optimal solution τ_* of (3). This is crucial for the final step of the proof.

Proving $\tilde{\mathbf{w}} \in \mathcal{S}$. Recall the definition in (10). We prove that $\frac{1}{n}\sum_{i=1}^{n} \mathbb{1}_{\{\tilde{\mathbf{w}}_i \leq -1\}} \xrightarrow{P} Q(1/\tau_*)$. From (11), $\mathbb{1}_{\{\tilde{\mathbf{w}}_i \leq -1\}} = \mathbb{1}_{\{\mathbf{h}_i \leq -\frac{\|\mathbf{g}\|}{\sqrt{n}\sqrt{\delta}\tau}\}}$. Recall, $\|\mathbf{g}\|/\sqrt{n} \xrightarrow{P} \sqrt{\delta}$ and $\tau \xrightarrow{P} \tau_*$. Conditioning on those high probability events it can be shown that $\frac{1}{n}\sum_{i=1}^{n} \mathbb{1}_{\{\mathbf{h}_i \leq -\frac{\|\mathbf{g}\|}{\sqrt{n}\sqrt{\delta}\tau}\}} \xrightarrow{P} \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}_{\{\mathbf{h}_i \leq -\frac{1}{\tau_*}\}} \xrightarrow{P} Q(\frac{1}{\tau_*})$.

4. DISCUSSION AND CONCLUSION

In this paper we have used the CGMT framework of [14] to precisely compute the P_e of the box relaxation method (BRO) to recover BPSK signals in MIMO systems. At high SNR we obtain $P_e = Q(\sqrt{(\delta - 1/2)SNR})$, compared to the Mathced filter bound (MFB), $Q(\sqrt{\delta SNR})$. As the interested reader may observe and expect, similar results can be achieved for higher order constellations (m-PAM, m-QAM, m-PSK, etc.). However, we shall leave the detailed calculations for another occasion.

In the proof outline, we made use of the set

$$\mathcal{S} = \{\mathbf{v} : \left| \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}_{(\mathbf{v}_i \leq -1)} - Q(\frac{1}{\tau^*}) \right| < \epsilon\},$$

to establish that the (AO) and (PO) have the same expected BER. A study of our analysis of the (AO) reveals that error events for each of the bits in the (AO) are *iid*. This means that if, for constant k, we define the set:

$$\mathcal{S}_k^* = \{\mathbf{v} : \left| \frac{1}{\binom{n}{k}} \sum_{\substack{T\in\{1,\ldots,n\} \\ |T|=k}} \mathbb{1}_{(\mathbf{v}_{i_1} \leq -1, \ldots, \mathbf{v}_{i_k} \leq -1)} - Q^k(\frac{1}{\tau^*}) \right| < \epsilon\}$$

then $\lim_{n\to\infty} \mathbb{P}\{\mathbf{w}_\phi \in S_k^*\} = 1$. By Thm. 2.1, this implies $\lim_{n\to\infty} \mathbb{P}\{\mathbf{w}_\Phi \in S_k^*\} = 1$, which means that error events for any fixed k bits in the (PO) are also iid. This fact has significant consequences. For example, it implies that, when a block of data is in error, only a few of its bits are. This means that the output of the (BRO) can be used by various local methods to further reduce the BER. We shall explain this in future work.

References

[1] H. Q. Ngo, E. G. Larsson, and T. L. Marzetta, "Energy and spectral efficiency of very large multiuser mimo systems," *Communications, IEEE Transactions on*, vol. 61, no. 4, pp. 1436–1449, 2013.

[2] C.-K. Wen, J.-C. Chen, K.-K. Wong, and P. Ting, "Message passing algorithm for distributed downlink regularized zero-forcing beamforming with cooperative base stations," *Wireless Communications, IEEE Transactions on*, vol. 13, no. 5, pp. 2920–2930, 2014.

[3] T. L. Narasimhan and A. Chockalingam, "Channel hardening-exploiting message passing (chemp) receiver in large-scale mimo systems," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 8, no. 5, pp. 847–860, 2014.

[4] J. Charles, G. Ramina, M. Arian, and S. Christoph, "Optimality of large mimo detection via approximate message passing," in *Information Theory (ISIT), 2015 IEEE International Symposium on*. IEEE, 2015.

[5] M. Grötschel, L. Lovász, and A. Schrijver, *Geometric algorithms and combinatorial optimization*. Springer Science & Business Media, 2012, vol. 2.

[6] G. J. Foschini, "Layered space-time architecture for wireless communication in a fading environment when using multi-element antennas," *Bell labs technical journal*, vol. 1, no. 2, pp. 41–59, 1996.

[7] B. Hassibi and H. Vikalo, "On the sphere-decoding algorithm i. expected complexity," *Signal Processing, IEEE Transactions on*, vol. 53, no. 8, pp. 2806–2818, 2005.

[8] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2009.

[9] S. Verdu, *Multiuser detection*. Cambridge university press, 1998.

[10] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, "The convex geometry of linear inverse problems," *Foundations of Computational Mathematics*, vol. 12, no. 6, pp. 805–849, 2012.

[11] M. Stojnic, "A framework to characterize performance of lasso algorithms," *arXiv preprint arXiv:1303.7291*, 2013.

[12] S. Oymak, C. Thrampoulidis, and B. Hassibi, "The squared-error of generalized lasso: A precise analysis," *arXiv preprint arXiv:1311.0830*, 2013.

[13] C. Thrampoulidis, A. Panahi, D. Guo, and B. Hassibi, "Precise error analysis of the lasso," in *40th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) 2015, arxivPreprint arXiv:1502.04977*, 2015.

[14] C. Thrampoulidis, S. Oymak, and B. Hassibi, "Regularized linear regression: A precise analysis of the estimation error," in *Proceedings of The 28th Conference on Learning Theory*, 2015, pp. 1683–1709.

[15] Y. Gordon, "Some inequalities for gaussian processes and applications," *Israel Journal of Mathematics*, vol. 50, no. 4, pp. 265–289, 1985.

[16] Y. Plan and R. Vershynin, "The generalized lasso with non-linear observations," *arXiv preprint arXiv:1502.04071*, 2015.

[17] Y. Gordon, *On Milman's inequality and random subspaces which escape through a mesh in* $\mathbb{R}^n$. Springer, 1988.

[18] M. Ledoux and M. Talagrand, *Probability in Banach Spaces: isoperimetry and processes*. Springer, 1991, vol. 23.

[19] C. Thrampoulidis, E. Abbasi, and B. Hassibi, "Precise high-dimensional error analysis of regularized m-estimators under gaussian measurement designs," in *Communication, Control, and Computing (Allerton), 2015 53rd Annual Allerton Conference on*. IEEE, 2015.

[20] P. K. Andersen and R. D. Gill, "Cox's regression model for counting processes: a large sample study," *The annals of statistics*, pp. 1100–1120, 1982.

[21] W. K. Newey and D. McFadden, "Large sample estimation and hypothesis testing," *Handbook of econometrics*, vol. 4, pp. 2111–2245, 1994.